

Appl. No. 10/689,257
Amdt. dated August 15, 2006
Reply to Office Action of 15 May 2006

Amendments to the Drawings:

Formal drawings are submitted with this amendment. Additionally, FIG. 14A and FIG. 19E have been amended as shown by the annotated sheets marked in red to show the changes.

Attachment: Replacement sheets 1/10 – 10/10
Annotated sheets 6/10 and 9/10

REMARKS/ARGUMENTS

In response to paragraph 3 of the Office action, the title of the invention has been amended to recite “Method of Shifting Data Along Diagonals in a Group of Processing Elements to Transpose the Data.” If the examiner remains of the opinion that the title is not sufficiently descriptive, the examiner is invited to suggest an alternative title.

The informalities identified in paragraph 4 of the Office action have been corrected. Additionally, typographical errors in the G.B. patent application numbers in paragraphs [0049] and [0052] have been corrected, and in paragraph [0052] the G.B. patent application number has been updated with the corresponding patent number. In paragraph 57, a period has been added at the end of the third sentence of that paragraph.

In paragraph 5 of the Office action, the examiner has objected to the quality of Fig. 2, Fig. 5, and Figs. 6A – 10B. In response, a complete set of formal drawings is filed with this amendment. Each of the replacement sheets of formal drawings has been labeled “Replacement Sheet” in the page header.

In paragraph 6 of the Office action, the examiner has objected to Fig. 14A on the basis that “it appears that there should not be a diagonal arrow emanating from element 11.” In response, the diagonal has been eliminated from element 11 and is now shown as emanating from element 13. An annotated sheet 6/10 is enclosed showing the changes in red.

In paragraph 6 of the Office action, the examiner has stated that in Fig. 19E it is not clear why the counts are set to zero to seven instead of minus one to six. The examiner is correct and Fig. 19E has been amended. Additionally, an annotated sheet 9/10 is enclosed herewith illustrating the changes in red.

In paragraph 7 of the Office action, the drawings are objected to as failing to show every feature of the invention specified in the claims, specifically, “the examiner has been unable to find any figure showing shifting of data along as [sic a] diagonal. Instead, all the examiner is able to see is the shifting data to the east (Fig. 17A-H).” It is respectfully submitted that when Figs. 17A-H are compared to the original input data of Fig. 16A, it is seen that the data is shifted along diagonals.

For example, the original data of row 1 of Fig. 16A is presented in Fig. 17A. Thereafter, a first pair of shifts is performed which, in this example, is an east shift and a north shift of all of the data in the array. Examining row 2 of Fig. 16A, and comparing that to Fig. 17B, it is seen that the first row of processing elements (Fig. 17B) now has data from the second row of processing elements (Fig. 16A), but shifted to the east by one processing element. Thereafter, the process is repeated, i.e., pairs of east and north shifts are performed.

Looking again at Figs. 17A-17H, and focusing on the eighth or last processing element in the row, it is seen that that processing element originally contains h1, and then receives in sequence g2, f3, e4, d5, c6, b7, and a8. Now, if one looks at the original input data, it is seen that that data received by the last processing element of the first row as a result of seven pairs of east/north shifts is all the data in the diagonal of which the last processing element is a part. The same is true for every processing element in the array. Thus, it is respectfully submitted that Fig. 16A, viewed in combination with Figs. 17A-17H, clearly shows the feature of shifting data along diagonals of the processing elements until each processing element in the diagonal has received the data held by every other processing element in that diagonal. In view of the foregoing, it is respectfully requested that the objection to the drawings under 37 C.F.R. 1.83(a) be withdrawn.

In paragraphs 9 and 10 of the Office action, claims 1-10 and 26 are rejected under 35 U.S.C. § 112, first paragraph, as failing to comply with the enablement requirement. According to the Office, "[t]he claim(s) contain subject matter which was not described in the specification in such a way as to enable one skilled in the art to which it pertains, or with which it is most nearly connected, to make and/or use the invention." The examiner's attention is respectfully directed to the preceding discussion regarding how Figs. 16A and 17A-17H illustrate the shifting of data along diagonals of the processing elements until each of the processing elements in the diagonal has received the data held by every other processing element in that diagonal. The examiner's attention is also directed to paragraph [0071] which discusses Fig. 16A and 17A-17H, which is reproduced here for the examiner's convenience.

[0071] Turning to FIGs. 16A and 16B, FIG. 16B represents one possible transpose of the data of FIG. 16A. FIG. 17A illustrates the data appearing in the first row of the matrix of FIG. 16A while FIGs. 17B through 17H illustrate the data as it appears as a result of seven consecutive horizontal/vertical pairs of shifts, i.e., move data right (east) 1 PE, and move data up (north) one PE. In FIG. 17A, the

data "a1" appearing in position one represents that data in the first row of the transpose of FIG. 16B. So that data is selected as the final output. In FIG. 17B, after the first pair of east/north shifts, the data "a2" appearing in position two represents the data in the first row of the transpose of FIG. 16B. Similarly, in FIGs. 17C through 17H the data "a3" to "a8" appears in positions three through eight, respectively. Using the expression $(X+Y+1) \text{ MOD } (8)$, the local counters in each of the processing elements may be set to an initial count as illustrated in FIG. 18.

It is readily apparent, even to a person of less than ordinary skill in the art, that if data is moved to the right one PE and up one PE, in a series of consecutive horizontal/vertical pairs of shifts, that the data will be shifted along a diagonal.

Furthermore, Fig. 18 illustrates the count in each of the processing elements which enables that processing element to select, from amongst all of the data which it receives, the appropriate data to output so as to affect a transpose of the data. Examining the top row in Fig. 18, it is seen that the first processing element has a count of zero, corresponding to Fig. 17A, which indicates that the original data in that processing element is the data which that processing element should output to effect a transpose of the data in the array. Conversely, the eighth processing element in the first row has a count of seven, indicating that that processing element must wait for seven shifts, see Fig. 17B-17H, until it receives the data a8. The count of seven thus enables the eighth processing element in the first row to select as its output the data it receives after the seven pairs of shifts. It is respectfully submitted that paragraph [0071], Fig. 16A and Figs. 17A-17H illustrate how every element in a given diagonal can receive data held by every other element in that diagonal. Figure 18 clearly illustrates how each of the processing elements selects from amongst all of the data which it has received, that data which is to be output to effect a transpose of the input data.

Finally, Figs. 19A-19L reiterate the foregoing by illustrating input data, output data, and counter values through the process of shifting data. See paragraphs [0072] – [0075].

In view of the foregoing, it is respectfully submitted that the specification and figures, as filed, clearly support the language in the claims which recites shifting the data along diagonals of the plurality of processing elements until each processing element in the diagonal has received the data held by every other processing element in that diagonal. Although the examiner states,

“it is not clear how a transpose can be performed while allowing each element to receive data held by every other element in the diagonal,” that is achieved, for example, by each element maintaining a count as shown in Fig. 18, so that that element may select the appropriate data for output, i.e., final output selected based on a processing element’s position. The examiner’s statement that “it is not clear how a transpose can be performed while allowing each element to receive data held by every other element in the diagonal” is powerful evidence of the novelty and non-obviousness of the present invention.

In paragraph 13 of the Office action, claim 1 is rejected under 35 U.S.C. § 112 because the limitation “the diagonal” in line 3 has insufficient antecedent basis. Claim 1 has been amended to provide the necessary antecedent basis. Accordingly, the rejection of claims 1-10 under 35 U.S.C. § 112, second paragraph, should now be overcome.

In paragraph 15 of the Office action, claims 1, 9-11, 19-20, and 26 are rejected under 35 U.S.C. § 102(b) as being anticipated by Hanounik et al. “Linear-time Matrix Transpose Algorithms Using Vector Register File with Diagonal Registers,” 2001 (hereinafter referred to as Hanounik). It is respectfully submitted that the examiner reads too much into Hanounik.

The method illustrated in Figure 1 of Hanounik appears substantially different from applicant’s method. Figure 1B shows the matrix after the first stage of transposition. It appears that the matrix is divided into four quadrants, and the data of the lower left-hand and upper right-hand quadrants are swapped. Thereafter, the matrix is again divided into quadrants, and the data in certain quadrants data swapped as shown. Finally, in Fig. 1D, it appears that data are swapped between certain processing elements as shown. The method in Fig. 1 of Hanounik thus involves the swapping of data quadrant by quadrant rather than the shifting of data along a plurality of diagonals until the processing elements in each of the plurality of diagonals has received the data held by every other processing element in that diagonal.

Secondly, the examiner states that “the element that originally holds value 12 should hold value 21 at the end of the transpose. When it finally does, value 21 is outputted as the final data in the transpose.” It appears, however, that the data in each processing element at the end of the process is the data to be output. There is no need for the processing elements to select from

amongst all of the data which they have received which data to output, i.e., no selecting of data as final output based on position.

It is respectfully submitted that the examiner has ignored how the process of Hanounik's Figure 1 is actually being performed, and has instead substituted a hindsight analysis for how the data could be moved. For the foregoing reasons, it is the applicant's position that Fig. 1 of Hanounik has not been demonstrated to anticipate the method set forth in independent claims 1 and 26.

In paragraph 24 of the Office action, claims 21-25 are rejected under 35 U.S.C. § 102(e) as being anticipated by U.S. Patent No. 6,968,447 to Apisdorf hereinafter "Apisdorf." The examiner points to Figs. 3, 4A, 4B, and 6 as well as column 13, lines 4-19 and column 14, lines 9-28. It is respectfully submitted that the examiner has misapprehended the teachings of Apisdorf.

The relevant sections cited by the examiner are reproduced below.

In an alternate embodiment, each processing element can include an end critical section signal counter that records receipt of end critical section signals. Generally, the counter can be initialized to some value, and each time the processing element processes a critical section, **the counter can be decremented. Eventually, the counter reaches zero, and the processing element suspends instruction execution before entering a new critical section. The processing element resumes processing when the counter is incremented upon receiving an end critical section signal from an adjacent, upstream, processing element.** If, for example, the counter is initialized at zero, processing elements downstream from a particular processing element can not process critical sections before receiving an end critical section signal. In one possible embodiment, state element 510 provides the physical implementation for an end critical section signal counter. (emphasis added)

This section of Apisdorf describes a combination of decrementing and incrementing. The count does not control which of a plurality of data values is to be selected as the final output.

The examiner, however, ends the citation to Apisdorf too soon. Continuing in column 13, beginning at line 20, Apisdorf discloses the following:

Consider, for example, a team including three processing elements. Since none of the processing elements can process a critical section without receiving an end critical section signal, one of the processors is designated as the "start" processing element. The end critical section signal counter of the start processing

element is incremented to at least one (e.g., by system software or a host controller), and the end critical section signal counters of the rest of the processing elements are zeroed. The team begins processing jobs in normal mode. Necessarily, the start processing element is the first processing element to encounter and process a critical section. The start processing element decrements its end critical section signal counter when it processes the critical section. After processing the critical section, the start processing element provides an end critical section signal to a second, downstream, processing element.

The second processing element increments its critical section end signal counter when the end critical section signal is received from the start processing element. The second processing element may now decrement the counter and process a critical section. Processing was suspended if the second processing element had encountered a critical section before receiving the end critical section signal. In the alternative, processing the next critical section will proceed uninterrupted if a critical section has not yet been encountered. As is apparent, the processing element downstream may not process a critical section until an end critical section signal is received from the second processing element. This advantageously ensures that no processing elements in the team process a critical section until an end critical section signal is received.

It is clear from the foregoing that the examiner is simply picking and choosing those portions of the reference favorable to his case, while ignoring other teachings of the reference which undercut the examiner's position. Although claim 21 has been amended to define x, y, C, and R, it is apparent that when Apisdorf is considered for all that it teaches, Apisdorf does not anticipate claim 21.

Applicant at this time has not submitted any arguments in support of the patentability of the dependent claims. It is believed that independent claims 1, 11, and 21 are now in condition for allowance such that all of the dependent claims which depend either directly or indirectly therefrom are also in condition for allowance.

Appl. No. 10/689,257
Amdt. dated August 15, 2006
Reply to Office Action of 15 May 2006

Applicant has made a diligent effort to place the instant application in condition for allowance. Accordingly, a Notice of Allowance for claims 1- 26 is respectfully requested. If the examiner is of the opinion that the instant application is in condition for disposition other than through allowance, the examiner is respectfully requested to contact applicant's attorney.



Respectfully submitted,





Edward L. Pencoske
Reg. No. 29,688
JONES DAY
One Mellon Center
500 Grant Street, Suite 3100
Pittsburgh, PA, USA, 15219
(412) 394-9531
(412) 394-7959 (Fax)
Attorneys for Applicant

Annotated Sheet

6/10



These Annotated
Sheet will be
submit to the



Annotated Sheet

9/10

	STEP 1 - INITIALIZE							
INPUT DATA	a1	b1	c1	d1	e1	f1	g1	h1
OUTPUT (TRANPOSED) DATA								
COUNT	0	1	2	3	4	5	6	7

Fig.19A

	STEP 2 - TEST (LOAD A1)							
INPUT DATA	a1	b1	c1	d1	e1	f1	g1	h1
OUTPUT (TRANPOSED) DATA	a1							
COUNT	0	1	2	3	4	5	6	7

Fig.19B

	STEP 3 - 1ST DECREMENT							
INPUT DATA	a1	b1	c1	d1	e1	f1	g1	h1
OUTPUT (TRANPOSED) DATA	a1							
COUNT	-1	0	1	2	3	4	5	6

Fig.19C

	STEP 4 - 1ST PAIR OF SHIFTS							
INPUT DATA	h2	a2	b2	c2	d2	e2	f2	g2
OUTPUT (TRANPOSED) DATA	a1							
COUNT	-1	0	1	2	3	4	5	6

Fig.19D

	STEP 5 - TEST (LOAD A2)							
INPUT DATA	h2	a2	b2	c2	d2	e2	f2	g2
OUTPUT (TRANPOSED) DATA	a1	a2						
COUNT	0	1	2	3	4	5	6	7
	-1	0	1	2	3	4	5	6

Fig.19E

	STEP 6 - 2ND DECREMENT							
INPUT DATA	h2	a2	b2	c2	d2	e2	f2	g2
OUTPUT (TRANPOSED) DATA	a1	a2						
COUNT	-2	-1	0	1	2	3	4	5

Fig.19F

	STEP 7 - 2ND PAIR OF SHIFTS							
INPUT DATA	g3	h3	a3	b3	c3	d3	e3	f3
OUTPUT (TRANPOSED) DATA	a1	a2						
COUNT	-2	-1	0	1	2	3	4	5

Fig.19G

•
•
•